

Efficient Closed Negotiation

Dmytro Tykhonov

Man-Machine Interaction Group

Outline

- Introduction to negotiation
- Analysis of negotiation dynamics
- Bayesian learning for opponent modeling and using an opponent model
- Human-Human and Human-Machine negotiation experiments
- Conclusions and Future Work

Introduction to Negotiation

Outline

- Toy-example
- Utility functions
- Negotiation outcome evaluation
- Negotiation strategies

Outline

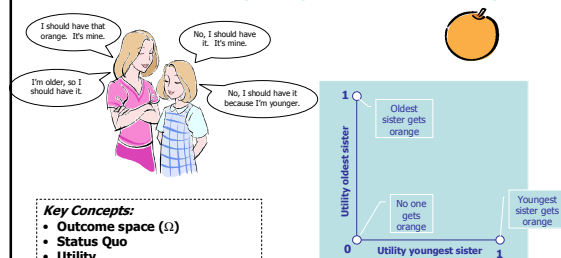
- Toy-example
- Utility functions
- Negotiation outcome evaluation
- Negotiation strategies

A Classic Tale: Two Sisters Arguing Over an Orange*




* The story about the sisters' conflict over the orange has been attributed to Mary Parker Follett. See Deborah M. Kolb, 1995, The Love for Three Oranges, Or: What Did We Miss about Ms. Follett in the Library? 11, Negotiation 1, 339. Also see: Roger Fisher & Danny Ertel, 1995, Getting Ready to Negotiate.

A Classic Tale: Two Sisters Arguing Over an Orange

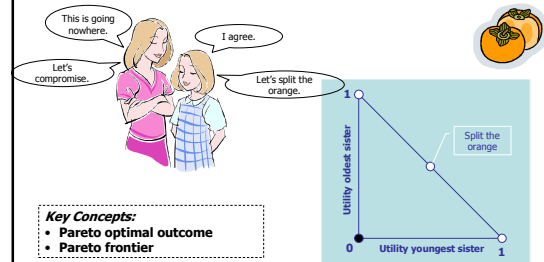


A Classic Tale: Two Sisters Arguing Over an Orange

The arguments are not decisive, but there is an alternative:




A Classic Tale: Two Sisters Arguing Over an Orange




A Classic Tale: Two Sisters Arguing Over an Orange

But why do both sisters want the orange?



A Classic Tale: Two Sisters Arguing Over an Orange

But why do both sisters want the orange?



Summary

- A negotiation is aimed at resolving a **conflict of interests**.
- Each negotiation is defined by its **outcome space**. Ideally, an agreement is reached that is Pareto optimal.
- During a negotiation, the outcome space is **explored**, and sometimes, may be **expanded** to find creative agreements based on the underlying interests.

Literature:

- H. Raiffa, 1982, *The Art and Science of Negotiation*, Belknap Press
- L.L. Thompson, 2005 (3rd ed.), *The Mind and Heart of the Negotiator*, Prentice Hall

Outline

- Toy-example
- Utility functions
- Negotiation outcome evaluation
- Aspects of negotiation
- Negotiation strategies

Where Do Utilities Come From?


- It is assumed that agents have preferences for outcomes, and thus are not completely indifferent about various outcomes.
- If preferences of an agent are both *rational* as well as *well-behaved* in a precise sense, then it can be shown that a utility function can be constructed that can be used for guiding the decision-making of that agent.
- Preferences are *complete* and *transitive*.

15

Predictable versus Unpredictable Preferences

Predictable preferences

- Most people prefer to have more rather than less money.




à Monotonic preference

à Logarithmic curve for money

Unpredictable preferences

- I like blue more than red, you like red more than blue.



à Varies from person to person

à Not quite as well-behaved

16

Utility: Determine the Value of a Bid

- Utility** is determined as a weighted average of the value of each of the issues.
- For example, if an accessory in a car negotiation such as a "cd-player" is very important, then the utility of the cd-player counts heavily in the total utility of a bid.

Attribuut	Waarde	Eval	Gewicht
cd_speler	Standard	0.7	0.9
airco	Standard	0.1	0.6
prijs	15000	0.6	0.9

$$\text{Utility} = (.9 \times .7 + .6 \times .1 + .9 \times .6) / (.9 + .6 + .9) \approx 0.51$$


$$\text{Utility} = (\text{Sum of weights} \times \text{values}) / (\text{Sum of weights})$$

17

Outline

- Toy-example
- Utility functions
- Negotiation outcome evaluation
- Negotiation strategies

How Good Did You Do?



20

Outline

- Toy-example
- Utility functions
- Negotiation outcome evaluation
- Negotiation strategies

Joint Exploration of Outcome Space

The Nash solution, and, more importantly, the Pareto frontier, cannot be computed when both parties only have incomplete preference information.

- Axiomatic approaches typically characterize possible outcomes of a negotiation, assuming **complete information**.
- In **closed negotiations**, negotiators at best have partial information about their opponent's preferences.
- For a good reason: Revealing information to an opponent may result in **exploitation**.
- Reaching an agreement requires negotiators with incomplete information to jointly **explore** the outcome space by **exchanging offers**, i.e. outcomes proposed to the other party.


24

Negotiation Decision functions

The complexity of negotiation with incomplete information does not allow for establishing game-theoretic optimal equilibrium strategy results. Instead negotiation heuristics are proposed in the literature.

Different families of negotiation heuristics:

- Time-dependent tactics
- Behaviour-dependent tactics



A tactic is used to determine a next offer.

25

Time Dependent Tactics

Time dependent tactics take deadlines and discounting of future rewards into account.

- A time dependent tactic determines the next offer to be proposed as a function $\alpha^a(t)$ of time:


$$x_{a \rightarrow b}^t = \begin{cases} \min^a + \alpha^a(t)(\max^a - \min^a) & \text{if } V^a \text{ decreasing} \\ \min^a + (1 - \alpha^a(t))(\max^a - \min^a) & \text{if } V^a \text{ increasing} \end{cases}$$

Various time-dependent functions are used:

- A polynomial function: $\alpha^a(t) = k^a + (1 - k^a) \left(\frac{\min(t, t_{\max}^a)}{t_{\max}^a} \right)^{\frac{1}{\beta}}$
- An exponential function: $\alpha^a(t) = e^{\left(\frac{\min(t, t_{\max}^a)}{t_{\max}^a} \right) \beta \ln k^a}$

26

Time Dependent Functions



27

Behavior Dependent Tactics

Behavior dependent tactics take the behavior of the opponent into account measured relative to the agent's own utility space, giving rise to tit-for-tat strategies.

- A behavior dependent tactic determines the next offer to be proposed as a function $\psi(t_n)$ of the last proposed offer t_n :

$$x_{a \rightarrow b}^{t_{n+1}} = \min(\max(\psi(t_n), \min^a), \max^a)$$

Various functions $\psi(t_n)$ are used:


- Average tit-for-tat: $\psi(t_n) = \frac{x_{b \rightarrow a}^{t_{n-2}}}{x_{b \rightarrow a}^{t_{n-1}}} x_{a \rightarrow b}^{t_{n-1}}$
- Relative tit-for-tat: $\psi(t_n) = \frac{x_{b \rightarrow a}^{t_{n-2}}}{x_{b \rightarrow a}^{t_{n-2}+2}} x_{a \rightarrow b}^{t_{n-1}}$

28

Smart Meta-Strategy

Using domain knowledge it may be possible to improve the efficiency of a strategy to the benefit of the agent itself as well as of its opponent.

- Key idea:** To better meet the demands of an opponent it is not always required to concede utility oneself.
- Propose alternative offer that is more *similar* to that of opponent's last offer with same own utility.
- Only concede if not possible.



- Result:** Issue trade-offs. *Example:* Propose (Medium Quality ↓, Low Price ↑, 10 days ⇒) instead of (High Quality, Medium Price, 10 days)
- See: Trade-Off Strategy (cf. literature)

30

Evaluating Strategies

A number of criteria are available to judge the quality of a strategy.

Process-Oriented:

- Cost-Effective: How many steps are required to reach agreement?
- Robustness: Is it possible to exploit the strategy?
- Negotiation move types: Is strategy sensitive to opponent?

Outcome-Oriented:

- Successful: Is an agreement reached?
- Efficiency: Is agreement Pareto optimal?
- Fairness: Is agreement close to Nash? Kalai-Smorodinsky?

31

Summary

- Due to incomplete information about opponent preferences, it is not possible to define "optimal" strategies in any precise sense.
- Various negotiation heuristics have been proposed, including **time** and **behavior dependent tactics**.
- Efficiency can be improved by using domain knowledge and proposing an offer on a utility iso-curve more similar to the opponent's last offer.

Literature:

- P. Faratin, C. Sierra, N.R. Jennings, 1997, *Negotiation Decision Functions for Autonomous Agents*, in: *Int. Journal of Robotics and Autonomous Systems*, 24(3-4), 159-182
- P. Faratin, C. Sierra, N.R. Jennings, 2002, *Using Similarity Criteria to Make Issue Trade-Offs in Automated Negotiations*, in: *Artificial Intelligence*, 142, 205-237
- K.V. Hindriks, C.M. Jonker, D. Tykhonov, 2007, *Analysis of Negotiation Dynamics*, in: *Cooperative Information Agents XI*, LNCS 4676, 27-35
- J.S. Rosenschein, G. Zlotkin, 1994, *Rules of Encounter*, MIT Press

32

Analysis of Negotiation Dynamics*

Outline

- Classes of negotiation moves and sensitivity measure;
- Experimental setup;
- Negotiation dynamics analysis of strategies and domains;
- Conclusions

33

* Koen Hindriks, Catholijn M. Jonker, and Dmytro Tykhonov; Negotiation Dynamics: Analysis, Concession Tactics, and Outcomes, *Proceedings of The 2007 IEEE/WIC/ACM International Conference on Intelligent Agent Technology (IAT'07)*, Fremont, USA, 2007.

Question

- To what extent negotiation moves (bids) affect the outcome of the negotiation?
 - What are the classes of negotiation moves?
 - What is the role of opponent model in negotiation dynamics?
 - How to measure sensitivity of a negotiation strategy to the opponent's preferences?

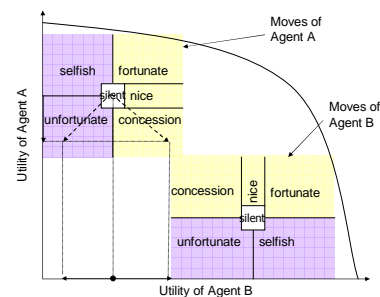
34

Outline

- Classes of negotiation moves and sensitivity measure;
- Experimental setup;
- Negotiation dynamics analysis of strategies and domains;

35

Classes of Moves



37

Sensitivity to Opponent Preferences

Assumption: a rational negotiator would try to make *fortunate, nice, or concession moves*.

$$\text{sensitivity}_a(t) = \frac{\%_{\text{Fortunate}}(t_a) + \%_{\text{Nice}}(t_a) + \%_{\text{Concession}}(t_a)}{\%_{\text{Selfish}}(t_a) + \%_{\text{Unfortunate}}(t_a) + \%_{\text{Silent}}(t_a)}$$

- In case no selfish, unfortunate or silent moves are made we stipulate that $\text{sensitivity}_a(t) = \infty$.
- If $\text{sensitivity}_a(t) < 1$, then an agent is more or less insensitive to opponent preferences;
- If $\text{sensitivity}_a(t) > 1$, then an agent is more or less sensitive to the opponent's preferences, with complete sensitivity for $\text{sensitivity}_a(t) = \infty$.

39

Outline

- Classes of negotiation moves and sensitivity measure;
- Experimental setup;
- Negotiation dynamics analysis of strategies and domains;

40

Experimental Setup - Tournament

- Three strategies:
 - ABMP*
 - Trade-Off**
 - Random-Walker
- Three domains:
 - Second hand car selling*
 - Service-oriented negotiation**
 - AMPO vs City***

41

(*) [Jonker and Treur, 2001]; (**) [Faratin *et al.*, 2003];
 (***) [Raiffa, 2002];

The Three Strategies

- ABMP
 - does not use any knowledge about opponent;
 - calculates concession size on every round of negotiation;
 - always make concession on every issue;
- Trade-off
 - uses domain knowledge;
 - tries to find bids on the same iso-level of own utility function that is closer to the current opponent's bid, makes concession of 0.05 if stuck;
 - uses opponent's bid to make trade-offs;
- Random-Walker (used as a benchmark)
 - Selects values of issues randomly
 - Proposes only those bids that have own utility > 0.6

42

The Three Domains

- Second hand car selling domain:
 - 5 issues (4 discrete issues and price issue),
 - only the buyer's preferences and the price issue are predictable
- Service-oriented negotiation (SON):
 - 4 continues issues;
 - all issues are predictable;
- AMPO vs City
 - 10 issues;
 - only 8 issues are predictable;

43


Outline

- Classes of negotiation moves and sensitivity measure;
- Experimental setup;
- Negotiation dynamics analysis of strategies and domains;

44


Dynamics of Negotiation Process

Trade-Off (City) vs Trade-Off strategy (AMPO)




Dynamics of Negotiation Process

Trade-Off (City) vs Random Walker strategy (AMPO)




Dynamics of Negotiation Process

Trade-Off (City) vs ABMP strategy (AMPO)



Dynamics of Negotiation Process

Random Walker (City) vs ABMP strategy (AMPO)



Analysis – Outcome Utility

- When considering overall utility the scores are: ABMP 0.72, Trade-Off 0.74, and Random Walker 0.69.
- Trade-Off:
 - Outperforms ABMP on the SON domain with complete information and on the AMPOvsCity domain;
 - Underperforms ABMP the second hand car domain due to wrong weights and unpredictable issues;
- ABMP:
 - Strong on the second hand car domain;
 - Underperforms on the SON domain;

Analysis – Sensitivity to Opponent Preferences

- Trade-Off:
 - perfectly sensitive to the opponent preferences on a domain with complete information (SON domain);
 - less sensitive on AMPOvsCity domain due to wrong weights;
 - low sensitivity on the second hand car domain due to wrong weights and unpredictable issues;
- ABMP:
 - very sensitive on a domain with complete information (SON domain);
 - rather stable sensitivity on the second hand car domain and the AMPOvsCity domain;

Conclusions and Future Work

- Conclusions
 - Want to negotiate efficiently? Know your partner!**
 - It is impossible to avoid unfortunate moves without sufficient domain knowledge or opponent knowledge.
 - In the analysis of negotiation strategies, not only the outcome of a negotiation is relevant, but also the bidding process itself is important.
 - The efficiency of the agreement does not correlate strongly with the sensitivity values of the strategies.
- Future work:
 - Apply proposed analytical methods to various negotiation strategies with opponent modeling;
 - Design negotiation testbed.

51

Opponent Modeling and Learning

Outline

- Opponent model structure and rationality assumption
- Bayesian learning for opponent modeling
- Using opponent model

52

Outline

- Opponent model structure and rationality assumption
- Bayesian learning for opponent modeling
- Using opponent model

Learning Opponent Preferences

Various features of an opponent's strategy can be learned, e.g. type of concession tactic, reservation value. We discuss learning opponent preferences.


- The strategies and tactics discussed so far are blind for opponent preferences.
- Various techniques have been proposed to learn features and construct an **opponent model**.
- Can information about **opponent preferences** be learned in single-instance negotiation?
 - Single-instance negotiation typical for e-commerce.
 - Hard because only bids exchanged can be used, no previous history can be assumed.

54

Learning: Structural Assumptions

In order to be able to learn opponent preferences we have to exploit certain structural features of negotiation problems.

- Linear additive utility functions: $u(b_i) = \sum_{i=1}^n w_i e_i(x_i \in b_i)$
- Sufficient to learn ranking of weights: $w_i = 2 \frac{r_i^j}{n(n+1)}$
- Shape of Evaluation functions (utility of issues):




Downhill,
Uphill, and
Triangular

55

Learning: Rationality Assumptions

In order to be able to learn opponent preferences we have to assume agents are (at least to some extent) rational.

- Opponent starts by proposing its best bid.
- Opponent uses a concession-based tactic
- Probability distributions are associated with possible tactics.



56

Outline

- Opponent model structure and rationality assumption
- Bayesian learning for opponent modeling
- Using opponent model

Bayesian Learning

Each (new) bid provides new information that is used to update the model of the opponent's preferences using Bayes' rule.

- Using Bayes' Rule the probabilities associated with hypotheses regarding weights and evaluation functions can be updated for every new bid b_i :


$$P(h_i | b_i) = \frac{P(h_i)P(b_i | h_i)}{\sum_{k=1}^n P(h_k)P(b_i | h_k)}$$

- An estimate of the utility an opponent associates with a bid can be computed using the computed probabilities:


$$\bar{u}(b_i) = \sum_{j=1}^{|U|} P(h_j) \sum_{i=1}^n w_i e_i(x_i \in b_i)$$

Learning: Scalability

Learn probabilities of the hypothesis about the evaluation functions and the weights independently of each other.



Experimental Results



- The Bayesian agent uses smart Meta-strategy
- AMPO vs City domain is used

Outline

- Opponent model structure and rationality assumption
- Bayesian learning for opponent modeling
- Using opponent model

Fundamental Problem of Negotiation

A dilemma faced by any negotiator is how to simultaneously achieve two typically conflicting goals: maximize own outcome and chance of a deal.

What is the optimal strategy to achieve the following:

- **Maximize own outcome.**
Requires outcome with high own utility, i.e. that is individually rational. *Rule of Thumb:* Start negotiation and check whether offer that is best for self is acceptable.
- **Maximize chance of reaching an agreement.**
Requires outcome with acceptable utility for opponent, i.e. resolving the conflict of interest.

Using an Opponent Model


An opponent model:

- enables choosing an offer that has utility x for the negotiator and maximizes the utility of the opponent (which helps to increase the chance of acceptance)
- enables to select an offer that increases the utility of both parties if previous offers have deviated from the Pareto frontier, which is typical. I.e. it allows making a **fortunate move**.

Defining an Opponent-based Strategy

Strategy using opponent model:

- Use a tit-for-tat tactic, i.e. match negotiation move that opponent just made.
- Propose Pareto-efficient offers, i.e. go "straight up" (agent A) to the Pareto frontier in order to maximize own utility.
- Many possible variants.



Summary

- It is possible to learn opponent preferences even in single-instance negotiations by making some rationality and structural assumptions.
- Efficiency can be improved by using learnt opponent preferences and by proposing an offer that increases utility of both parties.

Literature:

- D. Zeng, K. Sycara, 1998, *Bayesian Learning in Negotiation*, in: International Journal of Human-Computer Studies 48, 125-141
- R. Lin, S. Kraus, J. Wilkenfeld, J. Barry, 2006, *An Automated Agent for Bilateral Negotiation with Bounded Rational Agents with Incomplete Information*, ECAI 2006, 270-274
- K.V. Hindriks, D. Tykhonov, 2007, *Opponent Modelling in Automated Multi-Issue Negotiation Using Bayesian Learning*, Submitted.

Human-Human and Human-Machine experiments

Outline:

- Experimental setup
- Movie
- Results

Experimental Setup: Party Domain

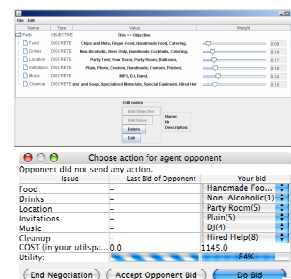
Issues:

- **Food** (Chips and nuts, finger-food, ...)
- **Drinks** (...)
- **Location** (Party tent, ...)
- **Invitations** (Plain, ...)
- **Music** (MP3, ...)
- **Cleanup** (Water and soap, ...)

Please note that there is a budget constraint of 1200 Euro. Bids that violate this constraint have a utility of 0!

Negotiation Session

- Prepare and fill in the preference profiles
- Negotiate with another human and sign an agreement (H-H)
- Negotiate with software agent (H-C)



Experimental Setup

- Incentive: the student's grade depends on the utility of the negotiation outcome
- In both session the students were negotiating with the same opponent (same preference profile) but were not aware of it
- The Bayesian Learning agent with a simple Tit-for-Tat strategy was used in the H-C session
- The guide and the software were made available to the students in advance

Movie

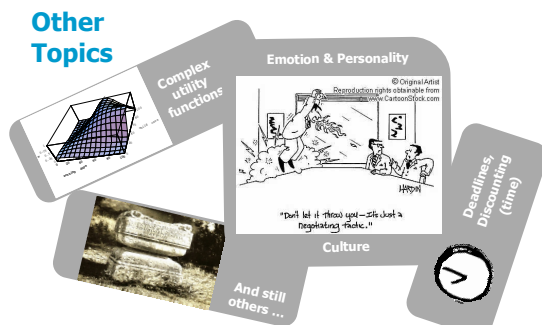
Results and Observations

- The students were well prepared for the experiment
- Open negotiation was preferred by the students
- The students tried to use calculators (Excel) to estimate utilities
- In 25 cases out of 30 (83%) the software agent had a better result than the human
- In 7 cases out of 30 (23%) the software agent had a better result than the human by 5%

Other Topics



Other Topics



Future Work – Pocket Negotiator (1)



Future Work – Pocket Negotiator (2)



Conclusion

What Does This Have to Do With Computing?

Negotiation Has Entered the Digital World:

- e-mail, e-commerce, ...

Artificial Intelligence:

- Engineering Heuristic Approaches for Machines
- Engineering Cognitive Skills used in Negotiation

The Future:

- Enhanced negotiations by using the Pocket Negotiator

